Exact Common Information

Gowtham Ramani Kumar
Electrical Engineering
Stanford University
Email: gowthamr@stanford.edu

Cheuk Ting Li
Electrical Engineering
Stanford University
Email: ctli@stanford.edu

Abbas El Gamal
Electrical Engineering
Stanford University
Email: abbas@stanford.edu

arXiv:1402.0062v1 [cs.IT] 1 Feb 2014

Abstract

This paper introduces the notion of exact common information, which is the minimum description length of the common randomness needed for the exact distributed generation of two correlated random variables (X, Y). We introduce the quantity $G(X;Y) = \min_{X \to W \to Y} H(W)$ as a natural bound on the exact common information and study its properties and computation. We then introduce the exact common information rate, which is the minimum description rate of the common randomness for the exact generation of a 2-DMS (X, Y). We give a multiletter characterization for it as the limit $\overline{G}(X;Y) = \lim_{n\to\infty}(1/n)G(X^n;Y^n)$. While in general $\overline{G}(X;Y)$ is greater than or equal to the Wyner common information, we show that they are equal for the Symmetric Binary Erasure Source. We do not know, however, if the exact common information rate has a single letter characterization in general.

I. Introduction

What is the common information between two correlated random variables or sources? This is a fundamental question in information theory with applications ranging from distributed generation of correlated sources [1] and secret keys [2] to joint source channel coding [3], among others. One of the most studied notions of common information is due to Wyner [1]. Let $(\mathcal{X} \times \mathcal{Y}, p(x, y))$ be a 2-DMS (or correlated sources (X, Y) in short). The Wyner common information $J(X;Y)$ between the sources X and Y is the minimum common randomness rate needed to generate (X, Y) with asymptotically vanishing total variation. Wyner established the single-letter characterization

$$J(X;Y) = \min_{W:\, X \to W \to Y} I(W; X, Y).$$

In this paper we introduce the notion of *exact common information*, which is closely related in its operational definition to the Wyner common information. While the Wyner setup assumes block codes and *approximate* generation of the 2-DMS (X, Y), our setting assumes variable length codes and *exact* generation of (X, Y). As such, the relationship between our setup and Wyner's is akin to that between the zero-error and the lossless source coding problems. In the source coding problem the entropy of the source is the limit on both the zero-error and the lossless compression. Is the limit on the exact common information rate the same as the Wyner common information? We show that they are the same for the Symmetric Binary Erasure Source (SBES) as defined in Section II. We do not, however, know if they are equal in general.

The rest of this paper is organized as follows. In the next section we introduce the exact distributed generation problem and define the exact common information. We introduce the "common-entropy" quantity $G(X;Y) = \min_{X \to W \to Y} H(W)$ as a natural bound on the exact common information and study some of its properties. In Section III, we define the exact common information rate for a 2-DMS. We show that it is equal to the limit $\overline{G}(X;Y) = (1/n)G(X^n;Y^n)$ and that it is in general greater than or equal to the Wyner common information. One of the main results in this paper is to show that $\overline{G}(X;Y) = J(X;Y)$ for the SBES. A consequence of this result is that the quantity $G(X^k;Y^k)$ can be strictly smaller than $kG(X;Y)$, that is, the per-letter common entropy can be reduced by increasing the dimension. We then introduce the notion of approximate common information rate, which relaxes the condition of exact generation to asymptotically vanishing total variation and show that it is equal to the Wyner common information. As computing the quantity $G(X;Y)$ involves solving a non-convex optimization problem, in Section V we present cardinality bounds on W and use them to find an explicit expression for $G(X;Y)$ when X and Y are binary. Due to space limitation, we do not include many of the proofs. We also mention a connection to the matrix factorization problem in machine learning that would be interesting to explore further.

II. Definitions and Properties

Consider the distributed generation setup depicted in Figure 1. Alice and Bob both have access to common randomness W. Alice uses W and her own local randomness to generate X and Bob uses W and his own local randomness to generate Y such that $(X, Y) \sim p_{X,Y}(x, y)$. We wish to find the limit on the least amount of common randomness needed to generate (X, Y) exactly.

More formally, we define a *simulation code* (W, R) for this setup to consist of

This work was partially supported by Air Force grant FA9550-10-1-0124.

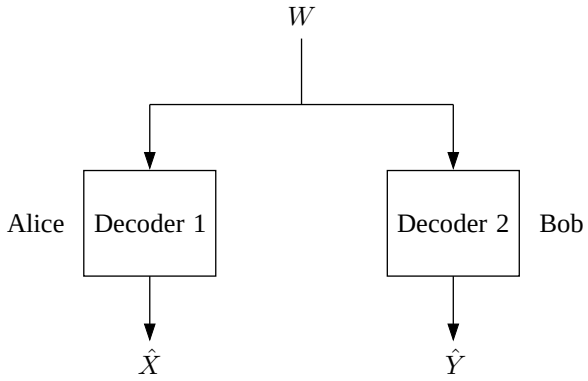

Fig. 1. Setting for distributed generation of correlated random variables. For exact generation, $(\hat{X}, \hat{Y}) \sim p_{X,Y}(x, y)$.

- A common random variable $W \sim p_W(w)$. As a measure of the amount of common randomness, we use the per-letter *minimum expected codeword length* R over the set of all variable length *prefix-free* zero-error binary codes $\mathcal{C} \subset \{0,1\}^*$ for W, i.e., $R = \min_{\mathcal{C}} \mathsf{E}(L)$, where L is the codeword length of the code $\mathcal{C}$ for W.
- A stochastic decoder $p_{\hat{X}|W}(x|w)$ for Alice and a stochastic decoder $p_{\hat{Y}|W}(y|w)$ for Bob such that $\hat{X}$ and $\hat{Y}$ are conditionally independent given W.

The random variable pair (X, Y) is said to be exactly generated by the simulation code (W, R) if $p_{\hat{X},\hat{Y}}(x, y) = p_{X,Y}(x, y)$. We wish to find the *exact common information* R^* between the sources X and Y, which is the infimum over all rates R such that the random variable pair (X, Y) can be exactly generated.

Define the following quantity, which can be interpreted as the "common entropy" between X and Y,

$$G(X;Y) = \min_{W:\, X \to W \to Y} H(W). \tag{1}$$

Remark: We can use $\min$ instead of $\inf$ in the definition of $G(X;Y)$ because the cardinality of W is bounded as we will see in Proposition 5, hence the optimization for computing $G(X;Y)$ is over a closed set.

Following the proof of Shannon's zero-error compression theorem, we can readily show the following.

Proposition 1.

$$G(X;Y) \le R^* < G(X;Y) + 1.$$

Computing $G(X;Y)$ is in general quite difficult (see Section V). In some special cases, we can find an explicit expression for it.

Example 1 The Symmetric Binary Erasure Source (SBES) (X, Y) is defined by

$$X \sim \mathrm{Bern}(1/2),$$
$$Y = \begin{cases} X & \text{w.p. } 1-p, \\ e & \text{w.p. } p, \end{cases}$$

where p is the *erasure probability* for the source. It can be shown that for the SBES,

$$G(X;Y) = \min\{1, H(p) + 1 - p\}.$$

Note that the Wyner common information for this source is [4]

$$J(X;Y) = \begin{cases} 1 & \text{if } p \le 0.5, \\ H(p) & \text{if } p > 0.5. \end{cases}$$

In the following we present some basic properties of $G(X;Y)$.

A. Properties of $G(X;Y)$

1) $G(X;Y) \ge 0$ with equality if and only if X and Y are independent.

   Proof: First assume $G(X;Y) = 0$. Suppose W achieves $G(X;Y)$. Then $X \to W \to Y$ and $H(W) = 0$. Thus $W = \phi$, constant. Hence X, Y are independent.

   To show the converse, if X, Y are independent, $X \to \phi \to Y$ and $G(X;Y) \le H(\phi) = 0$. Thus $G(X;Y) = 0$. ■

2) $G(X;Y) \ge J(X;Y)$.

Proof:

$$\begin{aligned} G(X;Y) &= \min_{W:X\to W\to Y} H(W) \\ &= H(W*) \\ &\geq I(W*;X,Y) \\ &\geq \min_{W:X\to W\to Y} I(W;X,Y) \\ &= J(X;Y). \end{aligned}$$

■

3) *Data-processing Inequality:* If $U \to X \to Y$ forms a Markov chain, then $G(U;Y) \leq G(X;Y)$.
   Proof: Let W achieve $G(X;Y)$. Then $U \to X \to W \to Y$ forms a Markov chain. Hence, $G(U;Y) \leq H(W) = G(X;Y)$. ■
4) Define $G(X;Y|Z) = \sum_{z\in\mathcal{Z}} p_Z(z) G(X;Y|Z=z)$. Then $G(X;Y) \leq H(Z) + G(X;Y|Z)$.
   Proof: For each $Z = z$, choose W_z as the random variable that achieves $G(X;Y|Z=z)$. Then $X \to (Z, W_Z) \to Y$ forms a Markov chain. Therefore,

$$\begin{aligned} G(X;Y) &\leq H(Z, W_Z) \\ &= H(Z) + H(W_Z|Z) \\ &= H(Z) + \sum_z p_Z(z) H(W_z) \\ &= H(Z) + G(X;Y|Z). \end{aligned}$$

■

5) If there exist functions $f(X)$ and $g(Y)$ such that $Z = f(X) = g(Y)$, then $G(X;Y) = H(Z) + G(X;Y|Z)$.
   Proof: First observe that if Z satisfies the condition $Z = f(X) = g(Y)$ and W satisfies the Markov condition $X \to W \to Y$, then Z is a function of W. To see why, note that

$$p_{X,Y}(x,y) = \sum_{w\in\mathcal{W}} p_W(w) p_{X|W}(x|w) p_{Y|W}(y|w).$$

   Therefore, if $p_{X,Y}(x,y) = 0$ and $p_{X|W}(x|w) > 0$, then $p_{Y|W}(y|w) = 0$.
   Now we will show that if $p_{X|W}(x_1|w) > 0$ and $p_{X|W}(x_2|w) > 0$, then $f(x_1) = f(x_2)$. If not, for any y such that $g(y) \neq f(x_1)$, $p_{X,Y}(x_1,y) = 0$, therefore $p_{X|W}(x_1|w) > 0 \implies p_{Y|W}(y|w) = 0$. Similarly for any y such that $g(y) \neq f(x_2)$, $p_{X,Y}(x_2,y) = 0$, therefore, $p_{X|W}(x_2|w) > 0 \implies p_{Y|W}(y|w) = 0$. As a consequence, for any y, $p_{Y|W}(y|w) = 0$, a contradiction.
   Thus, for any w the set $\{f(x) : p_{X|W}(x|w) > 0\}$ has exactly one element. If we define $h(w)$ as that unique element, then $f(X) = h(W)$. Therefore, $Z = f(X) = g(Y) = h(W)$.
   To complete the proof, let W achieve $G(X;Y)$. Then,

$$\begin{aligned} G(X;Y) &= H(W) = H(W,Z) \\ &= H(Z) + H(W|Z) \\ &= H(Z) + \sum_z p_Z(z) H(W|Z=z) \\ &\geq H(Z) + \sum_z p_Z(z) G(X;Y|Z=z) \\ &= H(Z) + G(X;Y|Z). \end{aligned}$$

   This, in combination with Property 4, completes the proof. ■
6) Let $T(X)$ be a sufficient statistic of X with respect to Y ([5], pg. 305). Then $G(X;Y) = G(T(X);Y)$. Further, if W achieves $G(X;Y)$, we have $H(W) \leq H(T(X))$. Thus a noisy description of X via W may potentially have a smaller entropy than the minimal sufficient statistic, which is a deterministic description.
   Proof: Observe that both Markov chains $T(X) \to X \to Y$ and $X \to T(X) \to Y$ hold. Hence by the data-processing inequality, $G(X;Y) = G(T(X);Y)$. Now, since W achieves $G(X;Y)$ and $T(X) \to T(X) \to Y$,

$$H(W) = G(X;Y) = G(T(X);Y) \leq H(T(X)).$$

■

III. Exact Common Information Rate

The distributed generation setup in Figure 1 can be readily extended to the n-letter setting in which Alice wishes to generate X^n from common randomness W_n and her local randomness and Bob wishes to generate Y^n from W_n and his local randomness such that $p_{\hat{X}^n,\hat{Y}^n}(x^n,y^n) \sim \prod_{i=1}^n p_{X,Y}(x_i,y_i)$. We define a simulation code (W_n,R,n) for this setup in the same manner as for the one-shot case.

We say that Alice and Bob can exactly generate the 2-DMS (X,Y) at rate R if for some $n \geq 1$, there exists a (W_n,R,n) simulation code that exactly generates (X^n,Y^n) (since we assume prefix-free codes for W_n, we can simulate for arbitrarily large lengths via concatenation of successive codewords). We wish to find the *exact common information rate* R^* between the sources X and Y, which is the infimum over all rates R such that the 2-DMS (X,Y) can be exactly generated.

Define the "joint common entropy"

$$G(X^n;Y^n) = \min_{W_n:\, X^n \to W_n \to Y^n} H(W_n). \tag{2}$$

It can be readily shown that $\lim_{n\to\infty}(1/n)G(X^n;Y^n) = \inf_{n\in\mathbb{N}}(1/n)G(X^n;Y^n)$. Hence, we can define the limiting quantity

$$\overline{G}(X;Y) = \lim_{n\to\infty} \frac{1}{n} G(X^n;Y^n).$$

We are now are ready to establish the following multiletter characterization for the exact common information rate.

Proposition 2 (Multiletter Characterization of R^*)**.** *The exact common information rate between the components X and Y of a 2-DMS (X,Y) is*

$$R^* = \overline{G}(X;Y).$$

Proof: Achievability: Suppose $R > \overline{G}(X;Y)$. We will show that the rate R is achievable.

Since $\overline{G}(X;Y) = \lim_{n\to\infty}(1/n)G(X^n;Y^n)$, $R \geq (1/n)(G(X^n;Y^n)+1)$ for n large enough. By the achievability part of Shannon's zero-error source coding theorem, it is possible to exactly generate (X^n,Y^n) at rate at most $(1/n)\left(G(X^n;Y^n)+1\right)$. Hence rate R is achievable and thus $R^* \leq \overline{G}(X;Y)$.

Converse: Now suppose a rate R is achievable. Then there exists a (W_n,R,n)- simulation code that exactly generates (X^n,Y^n). Therefore, by the converse for Shannon's zero-error source coding theorem, $R \geq (1/n)G(X^n;Y^n)$. Since $(1/n)G(X^n;Y^n) \geq \overline{G}(X;Y)$, we conclude that $R^* \geq \overline{G}(X;Y)$. ∎

As expected the exact common information rate is greater than or equal to the Wyner common information.

Proposition 3.

$$\overline{G}(X;Y) \geq J(X;Y).$$

The proof of this result is in Appendix C.

In the following section, we show that they are equal for the SBES in Example 1. We do not know if this is the case in general, however.

A. Exact Common Information of The SBES

We will need the following result regarding computing the Wyner common information for the SBES.

Lemma 1. *To compute $J(X;Y)$ for the SBES, it suffices to consider W of the form*

$$W = \begin{cases} X & \text{w.p. } 1-p_1, \\ e & \text{w.p. } p_1, \end{cases}$$

and

$$Y = \begin{cases} W & \text{w.p. } 1-p_2, \\ e & \text{w.p. } p_2, \end{cases}$$

where p_1,p_2 satisfy $p_1+p_2-p_1p_2 = p$, the erasure probability of the SBES.

The proof follows by [4], Appendix A.

We now present the main result on exact common information rate in this paper.

Theorem 1. *If (X,Y) is an SBES, then $\overline{G}(X;Y) = J(X;Y)$.*

Proof: In general $\overline{G}(X;Y) \geq J(X;Y)$. We will now provide an achievability scheme to show that for SBES, $\overline{G}(X;Y) \leq J(X;Y)$.

Choose a W as defined in Lemma 1 and define

$$\tilde{W} = \begin{cases} d & \text{if } W \in \{0,1\}, \\ e & \text{if } W = e, \end{cases}$$

$$\tilde{Y} = \begin{cases} d & \text{if } Y \in \{0,1\}, \\ e & \text{if } Y = e. \end{cases}$$

Note that $\tilde{Y}^n$, denoting the location of the erasures, is i.i.d. $\mathrm{Bern}(p)$ (with $1 \leftarrow e$, $0 \leftarrow d$) and independent of X^n. Furthermore, Y^n is a function of X^n and $\tilde{Y}^n$.

Codebook Generation: Generate a codebook $\mathcal{C}$ consisting of $2^{n(I(\tilde{Y};\tilde{W})+\epsilon)}$ sequences $\tilde{w}^n(m)$, $m \in [1 : 2^{n(I(\tilde{Y};\tilde{W})+\epsilon)}]$, that "covers" almost all the $\tilde{y}^n$ sequences except for a subset of small probability $\delta(\epsilon)$. By the covering lemma ([6], page 62), such a codebook exists for large enough n.

This lets us associate every covered sequence $\tilde{y}^n$ with a unique $\tilde{w}^n = \tilde{w}^n(\tilde{y}^n) \in \mathcal{C}$ such that $(\tilde{y}^n, \tilde{w}^n) \in \mathcal{T}_\epsilon^{(n)}$.

Define the random variable

$$\tilde{W}_n = \begin{cases} \tilde{w}^n(\tilde{y}^n) & \text{if } \tilde{y}^n \text{ is covered by } \mathcal{C}, \\ \tilde{y}^n & \text{if } \tilde{y}^n \text{ is not covered.} \end{cases} \tag{3}$$

Note that $\tilde{W}_n$ is a function of $\tilde{Y}^n$ and that the set of erasure coordinates in $\tilde{W}_n$ is a subset of those in $\tilde{Y}^n$.

Channel Simulation Scheme:

1) The central node generates $\tilde{W}_n$ defined in (3) and sends it to both encoders.
2) Encoder 2 (Bob) generates $\tilde{Y}^n \sim p_{\tilde{Y}^n|\tilde{W}_n}(\tilde{y}^n|\tilde{w}^n)$
3) The central node generates and sends to both encoders a message M comprising i.i.d. $\mathrm{Bern}(1/2)$ bits for only those coordinates i of X^n where $\tilde{W}_n(i) = d$. Thus $H(M) \le n(1 - p_1 + \delta(\epsilon))$.
4) Encoder 1 (Alice) generates the remaining bits of X^n not conveyed by M using local randomness. Then X^n is independent of $\tilde{W}_n, \tilde{Y}^n$ and is i.i.d. $\mathrm{Bern}(1/2)$.
5) Encoder 2 generates $Y^n = Y^n(\tilde{W}_n, X^n) = Y^n(\tilde{W}_n, M)$. He only needs the bits X_i such that $\tilde{Y}_i = d$, which are available via M.

To complete the proof, note that $X^n \to (\tilde{W}_n, M) \to Y^n$ forms a Markov chain. Therefore,

$$\begin{aligned} G(X^n;Y^n) &\le H(\tilde{W}_n, M) + 1 \le H(\tilde{W}_n) + H(M) + 1 \\ &\overset{(a)}{\le} H(\delta(\epsilon)) + (1-\delta(\epsilon))H(\tilde{W}_n|\tilde{W}_n \in \mathcal{C}) \\ &\quad + \delta(\epsilon)H(\tilde{W}_n|\tilde{W}_n \notin \mathcal{C}) + n(1 - p_1 + \delta(\epsilon)) + 1 \\ &\overset{(b)}{\le} H(\delta(\epsilon)) + (1-\delta(\epsilon))\log|\mathcal{C}| \\ &\qquad + \delta(\epsilon)\log|\tilde{\mathcal{Y}}^n| + n(1 - p_1 + \delta(\epsilon)) + 1 \\ &= n(I(\tilde{Y};\tilde{W}) + 1 - p_1 + \delta(\epsilon)) \\ &\overset{(c)}{=} n(I(W;X,Y) + \delta(\epsilon)), \end{aligned}$$

where (a) follows by the grouping lemma for entropy, since $\mathsf{P}\{\tilde{W}_n \notin \mathcal{C}\} = \mathsf{P}\{\tilde{Y}^n \text{ not covered}\} = \delta(\epsilon)$; (b) follows since entropy is upper bounded by $\log$ of the alphabet size; and (c) follows from the definition of mutual information and some algebraic manipulations.

If we let $n \to \infty$, we obtain $\overline{G}(X;Y) \le I(W;X,Y) + \delta(\epsilon)$ for any $\epsilon > 0$. Minimizing $I(W;X,Y)$ over all W from Lemma 1 completes the proof. ∎

Note that the single letter characterization of the Wyner common information for the 2-DMS $(X^k, Y^k) \sim \prod_{i=1}^k p_{X,Y}(x_i, y_i)$ is k times that of the 2-DMS $(X;Y)$, that is, $\min I(W;X^k,Y^k) = k \min I(W;X,Y)$. The same property holds for the Gács–Körner–Witsenhuesen common information [7], and for the mutual information. In the following we show that $G(X^k;Y^k)$ can be strictly smaller than $kG(X;Y)$. Hence, it is possible to realize gains in the "common entropy" when we increase the dimension.

By the fact that for the SBEC, $\overline{G}(X;Y) = H(p)$ for $p > 1/2$ and $G(X;Y) = \min\{1, H(p) + 1 - p\}$, there exists a p such that $\overline{G}(X;Y) < G(X;Y)$. Hence, we can show by contradiction that there exists a 2-DMS (X,Y) such that $G(X^2;Y^2) < 2G(X;Y)$. We can also give an explicit example of a 2-DMS (X,Y) such that $G(X^2;Y^2) < 2G(X;Y)$. Let

$$p_{X,Y} = \begin{bmatrix} 1/3 & 1/3 \\ 1/3 & 0 \end{bmatrix}.$$

Then, by Proposition 8, we have $G(X;Y) = H(1/3)$, where $H(p)$, $0 \le p \le 1$, is the binary entropy function. Now,

$$p_{X^2,Y^2} = \begin{bmatrix} 1/9 & 1/9 & 1/9 & 1/9 \\ 1/9 & 0 & 1/9 & 0 \\ 1/9 & 1/9 & 0 & 0 \\ 1/9 & 0 & 0 & 0 \end{bmatrix}.$$

If $X^2 \to W \to Y^2$, we can write

$$p_{X^2,Y^2}(x^2, y^2) = \sum_w p_W(w) p_{X^2|W}(x^2|w) p_{Y^2|W}(y^2|w).$$

Therefore, if we write

$$p_{X^2,Y^2} = \frac{4}{9} \begin{bmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}^t + \frac{3}{9} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 1/3 \\ 1/3 \\ 1/3 \end{bmatrix}^t + \frac{1}{9} \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}^t + \frac{1}{9} \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}^t.$$

we can identify $p_W(w) = \big[4/9, 3/9, 1/9, 1/9\big]$. It is easy to see $p_{W_1,W_2}(w_1, w_2) = \big[4/9, 2/9, 2/9, 1/9\big]$. Thus,

$$G(X^2;Y^2) = H(W) < H(W_1, W_2) = 2H(W_1) = 2G(X;Y).$$

B. Approximate common information rate

Consider the approximate distributed generation setting in which Alice and Bob wish to generate 2-DMS (X, Y) with vanishing total variation

$$\lim_{n\to\infty} \big| p_{\hat{X}^n,\hat{Y}^n}(x^n, y^n) - \prod_{i=1}^n p_{X,Y}(x_i, y_i) \big|_{\mathrm{TV}} = 0.$$

We define a (W_n, R, n)-simulation code for this setting in the same manner as for exact distributed generation. We define the *approximate common information rate* R^*_{TV} between the sources X and Y as the infimum over all rates R such that the 2-DMS (X, Y) can be approximately generated.

We can show that the approximate common information rate is equal to the Wyner common information.

Proposition 4.

$$R^*_{\mathrm{TV}} = J(X;Y).$$

Proof: Achievability: Achievability follows from Wyner's coding scheme [1]. Choose $W_n \sim \mathrm{Unif}[1 : 2^{nR}]$ and associate each $w_n \in \mathcal{W}_n$ with a codeword of fixed length $\ell(w_n) = \lceil nR \rceil$. Decoders 1 (Alice) and 2 (Bob) first decode W_n and then use Wyner's coding scheme to generate $\hat{X}^n, \hat{Y}^n$, respectively. Any rate $R > J(X;Y)$ is admissible and will guarantee the existence of a scheme such that $(\hat{X}^n, \hat{Y}^n)$ is close in total variation to (X^n, Y^n). Thus $R^*_{\mathrm{TV}} \le J(X;Y)$.

Converse: Suppose that for any $\epsilon > 0$, there exists a (W_n, R, n) simulation code that generates $(\hat{X}^n, \hat{Y}^n)$ whose pmf differs from that of (X^n, Y^n) by at most ϵ in total variation. Then we have

$$\begin{aligned}
nR &\ge H(W_n) \ge I(\hat{X}^n, \hat{Y}^n; W_n) \\
&= \sum_{q=1}^n I(\hat{X}_q, \hat{Y}_q; W | \hat{X}^{q-1}, \hat{Y}^{q-1}) \\
&= \sum_{q=1}^n I(\hat{X}_q, \hat{Y}_q; W, \hat{X}^{q-1}, \hat{Y}^{q-1}) \\
&\qquad\qquad - I(\hat{X}_q, \hat{Y}_q; \hat{X}^{q-1}, \hat{Y}^{q-1}) \\
&\overset{(a)}{\ge} \sum_{q=1}^n I(\hat{X}_q, \hat{Y}_q; W) - n\delta(\epsilon) \\
&= nI(\hat{X}_Q, \hat{Y}_Q; W | Q) - n\delta(\epsilon) \\
&= nI(\hat{X}_Q, \hat{Y}_Q; W, Q) - nI(\hat{X}_Q, \hat{Y}_Q; Q) - n\delta(\epsilon) \\
&\overset{(b)}{\ge} nI(\hat{X}_Q, \hat{Y}_Q; W, Q) - n\delta(\epsilon) \\
&\overset{(c)}{\ge} nJ(X;Y) - n\delta(\epsilon).
\end{aligned}$$

$(a), (b)$ follow from Lemma 20 and Lemma 21 respectively in [8] since the pmf of $(\hat{X}^n, \hat{Y}^n)$ differs from that of (X^n, Y^n) by at most ϵ in total variation; and (c) follows from the continuity of $J(X;Y)$ [1]. ■

Remark: Note that if we replace the total variation constraint in Proposition 4 by the stronger condition

$$p_{X^n,Y^n}(x^n, y^n) = (1-\epsilon)p_{\hat{X}^n,\hat{Y}^n}(x^n, y^n) + \epsilon r(x^n, y^n) \tag{4}$$

for some pmf $r(x^n, y^n)$ over $\mathcal{X}^n \times \mathcal{Y}^n$, the required approximate common information rate R^*_{SD} becomes equal to the exact common information $\overline{G}(X;Y)$. To show this, note that $R^*_{SD} \le \overline{G}(X;Y)$ is trivial because the exact distributed generation constraint is stronger than (4).

To show $R^*_{\mathrm{SD}} \ge \overline{G}(X;Y)$, start with any (W_n, R, n) simulation code that generates $(\hat{X}^n, \hat{Y}^n)$ satisfying (4). Let

$$W'_n = \begin{cases} W_n & \text{w.p. } 1-\epsilon, \\ (\bar{X}^n, \bar{Y}^n) \sim r(x^n, y^n) & \text{w.p. } \epsilon. \end{cases}$$

We construct a (W'_n, R', n) code that generates (X^n, Y^n) exactly and satisfies $R' \le R + \delta(\epsilon)$. If the decoders receive $W'_n = W_n$, they follow the original achievability scheme to generate $(\hat{X}^n, \hat{Y}^n)$ satisfying (4). If $W'_n = (\bar{X}^n, \bar{Y}^n)$, then the decoders simply output $\bar{X}^n$ and $\bar{Y}^n$, respectively. Now,

$$\begin{aligned} H(W'_n) &\le H(\epsilon) + (1-\epsilon)H(W_n) + \epsilon \log |\mathcal{X}|^n |\mathcal{Y}|^n \\ &= H(W_n) + n\delta(\epsilon). \end{aligned}$$

Therefore, $R' \le (1/n)(H(W'_n) + 1) = R + \delta(\epsilon) + 1/n = R + \delta(\epsilon)$ for n large enough. Thus $R^*_{SD} \ge \overline{G}(X;Y)$.

IV. Exact coordination capacity

In this section, we consider exact channel simulation, an extension of channel simulation with total variation constraint introduced in [9]. Consider the setup shown in Figure 2. Nature generates $X^n \sim \prod_{i=1}^n p_X(x_i)$ that is available to the encoder. Both encoder and decoder have access to common randomness W_n. The encoder sends a message $M(X^n, W)$ to the decoder. The decoder outputs $\hat{Y}^n$ using the message M, the common randomness W_n, and local randomness. We wish to characterize the trade-off between the amount of common randomness and the information rate.

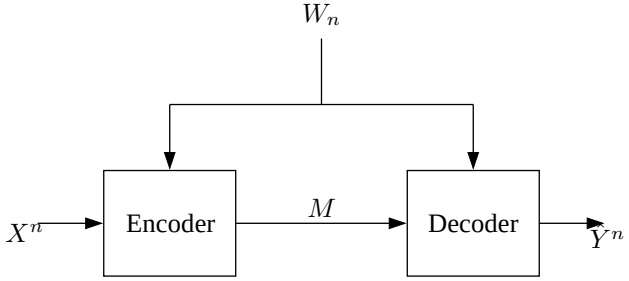

Fig. 2. Setting for exact channel simulation.

Formally, a (W_n, R, R_0, n) channel simulation code for this setup consists of

- a common random variable $W_n \sim p_{W_n}(w)$ independent of the source X^n. As a measure of the amount of common randomness, we use the per-letter *minimum expected codeword length* R_0 over the set of all variable length *prefix-free* zero-error binary codes $\mathcal{C}_0 \subset \{0,1\}^*$ for W_n, i.e., $R_0 = \min_{\mathcal{C}_0} \mathsf{E}(L)$, where L is the codeword length of the code $\mathcal{C}_0$ for W_n,
- an encoding function $M(X^n, W_n)$ that maps (X^n, W_n) into a random variable M. As a measure of the information rate, we use the per-letter *minimum expected codeword length* R over the set of all variable length *prefix-free* zero-error binary codes $\mathcal{C} \subset \{0,1\}^*$ for M, i.e., $R = \min_{\mathcal{C}} \mathsf{E}(L')$, where L' is the codeword length of the code $\mathcal{C}$ for M, and
- a stochastic decoder $p_{\hat{Y}^n|M,W_n}(y^n|m,w)$ that outputs $\hat{Y}^n$.

The channel simulation code is said to simulate the DMC $p_{Y|X}(y|x)$ exactly if $p_{\hat{Y}^n|X^n}(y^n|x^n) = \prod_{i=1}^n p_{Y|X}(y_i|x_i)$.

We wish to characterize the set of all achievable rates (R, R_0) for which the DMC $p_{Y|X}(y|x)$ can be simulated exactly.

We do not know the rate region for exact simulation of an arbitrary DMC. In the following, we show that for the erasure channel, it is equal to the rate region under total variation in [4].

Theorem 2. *When X is a binary symmetric source and $p_{Y|X}(y|x)$ is a binary erasure channel with erasure probability p, the rate region for exact channel simulation is the set of rate pairs (R, R_0) such that*

$$\begin{aligned} R &\geq r, \\ R + R_0 &\geq H(p) + r\big(1 - H((1-p)/r)\big), \end{aligned} \tag{5}$$

for some $1 - p \leq r \leq \min\{2(1-p), 1\}$.

Proof: The converse holds for total variation constraint and should therefore trivially hold for the more stringent exact channel simulation constraint. The achievability proof closely resembles that of Theorem 1. The central node still generates $\tilde{W}_n$, but the message M, instead of being generated at the central node, is now generated by the encoder (Alice). Thus, a rate pair (R, R_0) is achievable if

$$\begin{aligned} R &= (1/n)(H(M) + 1) + \delta(\epsilon) = 1 - p_1 + \delta(\epsilon), \\ R_0 &= (1/n)(H(\tilde{W}_n) + 1) + \delta(\epsilon) = I(\tilde{Y}; \tilde{W}) + \delta(\epsilon) = H(p) - (1-p_1)H(p_2) + \delta(\epsilon). \end{aligned} \tag{6}$$

for some p_1 and p_2 such that $p = p_1 + p_2 - p_1 p_2$. Letting $r = 1 - p_1$ shows that $R = r, R_0 = H(p) - rH((1-p)/r)$ is achievable. Finally note that both the encoder (Alice) and the central node can share the responsibility of generating and sending $\tilde{W}_n$ to the decoder (Bob). Thus an arbitrary fraction of R_0 can be removed and instead added to R. This shows the equivalence of (5) and (6).

■

V. Computing $G(X;Y)$

The optimization problem for determining $G(X;Y)$ is in general quite difficult, involving the minimization of a concave function over a complex Markovity constraint. In this section we provide some results on this optimization problem. We provide two bounds on the cardinality of W, establish two useful extremal lemmas, and use these results to analytically compute $G(X;Y)$ for binary alphabets. We then briefly discuss a connection to a problem in machine learning.

We first establish the following upper bound on cardinality.

Proposition 5. *To compute $G(X;Y)$ for a given pmf $p_{X,Y}(x,y)$, it suffices to consider W with cardinality $|\mathcal{W}| \leq |\mathcal{X}||\mathcal{Y}|$.*

The proof of this proposition is in Appendix D.

We now state an extremal lemma regarding the optimization problem for $G(X;Y)$ that will naturally lead to another cardinality bound.

Lemma 2. *Given $p_{X,Y}(x,y)$, let W attain $G(X;Y)$. Then for $w_1 \neq w_2$, the supports of $p_{Y|W}(\cdot|w_1)$ and $p_{Y|W}(\cdot|w_2)$ must be different.*

The proof of this lemma is in Appendix E.

Lemma 2 yields the following bound on the cardinality of W.

Proposition 6. *To compute $G(X;Y)$ for a given pmf $p_{XY}(x,y)$, it suffices to consider W with cardinality $|\mathcal{W}| \leq 2^{\min(|\mathcal{X}|,|\mathcal{Y}|)} - 1$.*

Proof: Suppose $|\mathcal{W}| > 2^{|\mathcal{Y}|} - 1$. Since there are only $2^{|\mathcal{Y}|} - 1$ non-empty subsets of $\mathcal{Y}$, by pigeon-hole principle, there exists $w_1 \neq w_2$ such that the supports of $p_{Y|W}(\cdot|w_1)$ and $p_{Y|W}(\cdot|w_2)$ are the same. This contradicts Lemma 2. Hence $|\mathcal{W}| \leq 2^{|\mathcal{Y}|} - 1$. By a symmetric argument, $|\mathcal{W}| \leq 2^{|\mathcal{X}|} - 1$. ■

The following shows that the bound in Proposition 6 is tight.

Example 2 Let (X, Y) be a SBES with $p = 0.1$. Since $p_{X,Y}(0,1) = p_{X,Y}(1,0) = 0$, the Markovity constraint $X \to W \to Y$ implies that the only W with $|\mathcal{W}| = 2$ is $W = X$; see [4], Appendix A. Hence, $G(X;Y) \leq H(X) = 1$. However, $H(Y) = H(0.1) + 0.1 < 1$. Thus, the optimal W^* that achieves $G(X;Y)$ requires $|\mathcal{W}^*| = 3$, making the bound in Proposition 6 tight.

The following is another extremal property of $G(X;Y)$.

Proposition 7. *Suppose W attains $G(X;Y)$. Consider a non-empty subset $\mathcal{W}' \subseteq \mathcal{W}$. Let (X', Y') be defined by the joint pmf*

$$p_{X',Y'}(x,y) = \sum_{w \in \mathcal{W}'} \frac{p_W(w)}{\sum_{w' \in \mathcal{W}'} p_W(w')} p_{X|W}(x|w) p_{Y|W}(y|w).$$

Then $H(X';Y') = H(W|W \in \mathcal{W}')$.

The proof of this proposition is in Appendix F.

We now use the above results to analytically compute $G(X;Y)$ for binary alphabets, i.e., when $|\mathcal{X}| = |\mathcal{Y}| = 2$.

Proposition 8. *Let* $X \sim \mathrm{Bern}(p)$ *and*

$$p_{Y|X} = \begin{bmatrix} \alpha & \beta \\ \bar{\alpha} & \bar{\beta} \end{bmatrix}$$

for some $\alpha, \beta \in [0,1], \bar{\alpha} = 1-\alpha, \bar{\beta} = 1-\beta$. *Let* W *achieve* $G(X;Y)$. *Then either*

$$p_{Y|W} = \begin{bmatrix} \alpha & 1 \\ \bar{\alpha} & 0 \end{bmatrix}, \ p_{W|X} = \begin{bmatrix} 1 & \bar{\beta}/\bar{\alpha} \\ 0 & 1-\bar{\beta}/\bar{\alpha} \end{bmatrix}, \text{ and}$$
$$W \sim \mathrm{Bern}\left(\bar{p}\left(1-\bar{\beta}/\bar{\alpha}\right)\right),$$

or

$$p_{Y|W} = \begin{bmatrix} 0 & \beta \\ 1 & \bar{\beta} \end{bmatrix}, \ p_{W|X} = \begin{bmatrix} 1-\alpha/\beta & 0 \\ \alpha/\beta & 1 \end{bmatrix}, \text{ and}$$
$$W \sim \mathrm{Bern}\left(p(1-\alpha/\beta)\right).$$

The proof of this proposition uses Lemma 2 as well as the cardinality bound $|\mathcal{W}| \leq 3$ derived from Proposition 6. It considers all possible cases for W and finally concludes that $|\mathcal{W}| = 2$ suffices. The detailed arguments can be found in Appendix G.

Remark (Relationship to machine learning): Computing $G(X;Y)$ is closely related to *positive matrix factorization*, which has applications in recommendation systems, e.g., [10]. In that problem, one wishes to factorize a matrix M with positive entries in the form $M = AB$, where A and B are both matrices with positive entries. Indeed, finding a Markov chain $X \to W \to Y$ for a fixed $p_{X,Y}$ is akin to factorizing $p_{Y|X} = p_{Y|W}p_{W|X}$ and numerical methods such as in [11] can be used. Rather than minimizing the number of factors $|\mathcal{W}|$ as is done in positive matrix factorization literature, it may be more meaningful for recommendation systems to minimize the entropy of the factors W. Computing $G(X;Y)$ for large alphabets appears to be very difficult, however.

VI. Conclusion

We introduced the notion of exact common information for correlated random variables (X,Y) and bounded it by the common entropy quantity $G(X;Y)$. For the exact generation of a 2-DMS, we established a multiletter characterization of the exact common information rate. While this multiletter characterization is in general greater than or equal to the Wyner common information, we showed that they are equal for the SBES. The main open question is whether the exact common information rate has a single letter characterization in general. Is it always equal to the Wyner common information? Is there an example 2-DMS for which the exact common information rate is strictly larger than the Wyner common information? It would also be interesting to further explore the application to machine learning.

VII. Acknowledgments

The authors are indebted to Young-Han Kim for comments that greatly improved the presentation of this paper.

Appendix A
Computing $G(X;Y)$ for SBES

As described in [4], Appendix A, for any W satisfying $X \to W \to Y$, each $w \in \mathcal{W}$ must fall in one of the three categories:

1) If $p_{X|W=w}p_{Y|W=w}$ is positive for $Y=0$, then $X=0$ with probability 1.
2) If $p_{X|W=w}p_{Y|W=w}$ is positive for $Y=1$, then $X=1$ with probability 1.
3) The last category is when $p_{X|W=w}p_{Y|W=w}$ is zero on $Y \in \{0,1\}$, i.e., $p_{Y|W}(e|w) = 1$.

Following the same reasoning in [4], Appendix A, we conclude that we need only one w in each category. Thus,

$$\begin{aligned} p_{X,Y} &= \begin{bmatrix} (1-p)/2 & 0 & p/2 \\ 0 & (1-p)/2 & p/2 \end{bmatrix} \\ &= \sum_{w \in \mathcal{W}} p_W(w) p_{X|W}(\cdot|w) p_{Y|W}(\cdot|w)^t \\ &= \begin{bmatrix} a_1 & 0 & a_2 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & b_1 & b_2 \end{bmatrix} + \begin{bmatrix} 0 & 0 & c_1 \\ 0 & 0 & c_2 \end{bmatrix}. \end{aligned}$$

Each term above corresponds to a category of w. Thus,

$$p_W = \begin{bmatrix} a_1 + a_2 & b_1 + b_2 & c_1 + c_2 \end{bmatrix}$$

and we must minimize $H(W)$ such that $a_1 = b_1 = (1-p)/2$, $a_2 + c_1 = b_2 + c_2 = p/2$, and $b_1, b_2, c_1, c_2 \geq 0$. This results in

$$G(X;Y) = \min_{0 \leq c_1, c_2 \leq p/2} H\left(1/2 - c_1, 1/2 - c_2, c_1 + c_2\right).$$

Since the objective is a concave function of c_1 and c_2, the minimizing (c_1, c_2) should be one of the 4 corner points $(0,0)$, $(p/2, p/2)$, $(0, p/2)$, $(p/2, 0)$. The last two points are symmetric. Thus

$$G(X;Y) = \min\{1, H(p) + 1 - p\}.$$

APPENDIX B
SUBADDITIVITY OF $G(X^n;Y^n)$

Suppose W_m achieves $G(X^m;Y^m)$ and W_n achieves $G(X^n;Y^n)$. Then, $X^m \to W_m \to Y^m$ and $X_{m+1}^{m+n} \to W_n \to Y_{m+1}^{m+n}$ form Markov chains, and so does $X^{m+n} \to (W_m, W_n) \to Y^{m+n}$. Therefore,

$$\begin{aligned} G(X^{m+n};Y^{m+n}) &\leq H(W_m, W_n) \\ &= H(W_m) + H(W_n) \\ &= G(X^m;Y^m) + G(X^n;Y^n). \end{aligned}$$

Thus the sequence $\{G(X^n;Y^n) : n \in \mathbb{N}\}$ is sub-additive. Hence, an appeal to Feteke's subadditivity lemma [12] shows that

$$\lim_{n\to\infty} \frac{1}{n} G(X^n;Y^n) = \inf_{n \in \mathbb{N}} (1/n) G(X^n;Y^n).$$

APPENDIX C
PROOF OF PROPOSITION 3

To show this consider

$$\begin{aligned} \overline{G}(X;Y) &= \lim_{n\to\infty} \min_{W: X^n \to W \to Y^n} \frac{1}{n} H(W) \\ &= \lim_{n\to\infty} \frac{1}{n} H(W_n^*) \\ &\geq \lim_{n\to\infty} \frac{1}{n} I(W_n^*; X^n, Y^n) \\ &\geq \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} I(W_n^*; X_i, Y_i) \\ &\geq \min_{W: X - W - Y} I(W; X, Y) \\ &= J(X;Y). \end{aligned}$$

APPENDIX D
PROOF OF PROPOSITION 5

We use the perturbation method (see Appendix C in [6]). The Markov Chain $X \to W \to Y$ is equivalent to

$$\begin{aligned} p(y|x,w) &= p(y|w), \text{ or} \\ \frac{p(x,y,w)}{\sum_{y'} p(x,y',w)} &= \frac{\sum_{x'} p(x',y,w)}{\sum_{x',y'} p(x',y',w)}. \end{aligned}$$

Further, $\sum_w p(x,y,w) = p_{X,Y}(w)$.

Let $p_\epsilon(x,y,w) = p(x,y,w)(1 + \epsilon\phi(w))$ be a perturbed pmf, where ϵ can be either positive or negative. We first observe that

$$\begin{aligned} p_\epsilon(y|x,w) &= \frac{p_\epsilon(x,y,w)}{\sum_{y'} p_\epsilon(x,y',w)} \\ &= \frac{p(x,y,w)}{\sum_{y'} p(x,y',w)} \\ &= \frac{\sum_{x'} p(x',y,w)}{\sum_{x',y'} p(x',y',w)} \\ &= \frac{\sum_{x'} p_\epsilon(x',y,w)}{\sum_{x',y'} p_\epsilon(x',y',w)} \\ &= p_\epsilon(y|w). \end{aligned}$$

Thus, $p_\epsilon(x,y,w)$ also satisfies the Markov chain $X \to W \to Y$.

Now, we also require that $p_\epsilon(x,y) = p_{X,Y}(x,y)$, i.e.,

$$\sum_w p(x,y,w)\phi(w) = 0 \text{ for all } (x,y).$$

The above equation represents $|\mathcal{X}||\mathcal{Y}|$ linear constraints in $\phi(w)$. Thus if $|\mathcal{W}| > |\mathcal{X}||\mathcal{Y}|$, we can find a perturbation $\phi(w) \neq 0$ such that $p_\epsilon(x,y,w)$ satisfies $\sum_w p_\epsilon(x,y,w) = p_{X,Y}(x,y)$ as well as the Markov chain $X \to W \to Y$.

If we choose ϵ small enough, we can ensure that $1 + \epsilon\phi(w) \geq 0$ for all w and thus $p_\epsilon(u,v,w)$ is a valid pmf.

Note that $p_\epsilon(w) = p(w)(1+\epsilon\phi(w))$ is a linear function of ϵ. Since the entropy is a concave function, $H(p_\epsilon(w))$ is a concave function of ϵ and is minimized by an extremal ϵ. Such an extremal ϵ makes $p_\epsilon(w) = 0$ for some w, thus reducing the cardinality $|\mathcal{W}|$ by 1 while also reducing $H(W)$. Thus we conclude for the minimizing $H(W)$, $|\mathcal{W}| \leq |\mathcal{X}||\mathcal{Y}|$.

APPENDIX E
PROOF OF LEMMA 2

We will prove by contradiction. Assume that Lemma 2 is false, i.e. $\exists w_1 \neq w_2$ such that the supports of $p_{Y|W}(\cdot|w_1)$ and $p_{Y|W}(\cdot|w_2)$ are the same for some W that achieves $H_1(X;Y)$.

Let $\mathcal{X} = \{1,2,\ldots,|\mathcal{X}|\}$, $\mathcal{Y} = \{1,2,\ldots,|\mathcal{Y}|\}$, etc. without loss of generality.

For a given w, $p_{Y|W}(\cdot|w_1)$ is a vector in $\mathrm{R}^{|\mathcal{Y}|}$ whose y-th element is the conditional probability $p_{Y|W}(y|w)$. The line segment joining $p_{Y|W}(\cdot|w_1)$ and $p_{Y|W}(\cdot|w_2)$ can be extended from both ends without crossing the boundary of the $|\mathcal{Y}|-1$ dimensional probability simplex as shown in Fig. 3.

Geometrically, the Markov chain $X \to W \to Y$ means that the vectors $\{p_{Y|X}(\cdot|x) : x \in \mathcal{X}\}$ lie in the convex hull of the vectors $\{p_{Y|W}(\cdot|w) : w \in \mathcal{W}\}$.

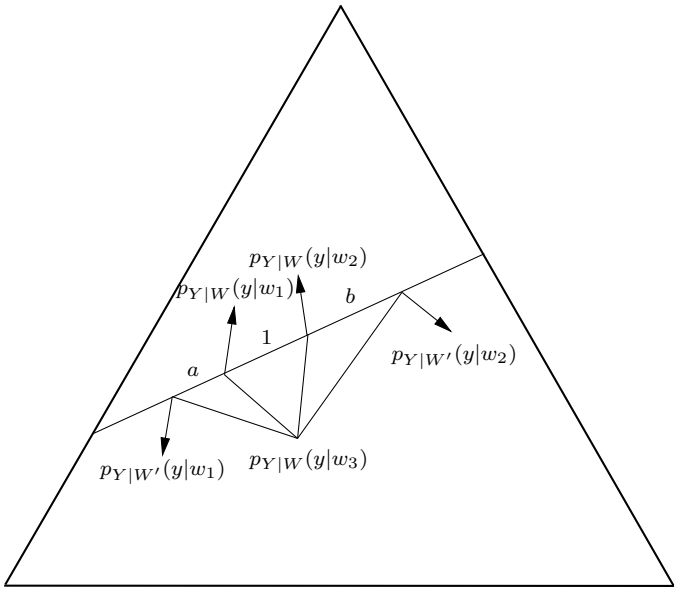

Fig. 3. Proof of Lemma 2.

We will construct W' such that $X \to W \to W' \to Y$ forms a Markov chain and $H(W') < H(W)$. Assume w.l.o.g. that the Euclidean distance between $p_{Y|W}(\cdot|w_1)$ and $p_{Y|W}(\cdot|w_2)$ in $\mathrm{R}^{|\mathcal{Y}|}$ is 1 unit. Extend the line segment in either direction by a, b units respectively so that it's new end-points are $p_{Y|W'}(\cdot|w_1), p_{Y|W'}(\cdot|w_2)$ (see Fig. 3). Then $X \to W \to W' \to Y$ forms a Markov chain, where W' is defined as follows:

$$p_{Y|W}(y|w) = \begin{cases} p_{Y|W'}(y|w) & \text{if } w \notin \{w_1, w_2\}, \\ \frac{b+1}{a+b+1} p_{Y|W'}(y|w_1) + \frac{a}{a+b+1} p_{Y|W'}(y|w_2) & \text{if } w = w_1, \\ \frac{b}{a+b+1} p_{Y|W'}(y|w_1) + \frac{a+1}{a+b+1} p_{Y|W'}(y|w_2) & \text{if } w = w_2. \end{cases}$$

and

$$p_{W'|W}(w'|w) = \begin{cases} \mathbf{1}(w'=w) & \text{if } w \notin \{w_1, w_2\}, \\ \frac{b+1}{a+b+1} & \text{if } (w', w) = (w_1, w_1), \\ \frac{a}{a+b+1} & \text{if } (w', w) = (w_2, w_1), \\ \frac{b}{a+b+1} & \text{if } (w', w) = (w_1, w_2), \\ \frac{a+1}{a+b+1} & \text{if } (w', w) = (w_2, w_2), \\ 0 & \text{otherwise.} \end{cases}$$

Hence we have

$$p_{W'}(\cdot) = \left[\frac{b+1}{a+b+1} p_W(w_1) + \frac{b}{a+b+1} p_W(w_2), \frac{a}{a+b+1} p_W(w_1) + \frac{a+1}{a+b+1} p_W(w_2), p_W(3), p_W(4), \dots \right],$$

and

$$H(W) - H(W') = \mathsf{P}\{W \in \{w_1, w_2\}\} \left[H(p) - H\left(\frac{b+p}{a+b+1} \right) \right], \tag{7}$$

where $p = \mathsf{P}\{W = w_1 | W \in \{w_1, w_2\}\}$ does not depend on a, b and $\bar{p} = 1 - p$. In (7), the term $(b+p)/(a+b+1)$ increases with b and decreases with a. Further, when $a = b = 0$, $(b+p)/(a+b+1) = p$. Hence, by perturbing either a or b, one can make the RHS of (7) positive. In that case $H(W') < H(W) = G(X;Y)$, a contradiction.

Appendix F
Proof of Proposition 7

Define the random variable $W' \in \mathcal{W}'$ with pmf $p_{W'}(w) = \frac{p_W(w)}{\sum_{w' \in \mathcal{W}'} p_W(w')}$. Then $X' \to W' \to Y'$ and $H(W') = H(W | W \in \mathcal{W}')$. Hence $H(X';Y') \le H(W') = H(W | W \in \mathcal{W}')$.

Now we show by contradiction that $H(X';Y') = H(W | W \in \mathcal{W}')$.

If $H(X';Y') < H(W | W \in \mathcal{W}')$, there exists an a $\tilde{W} \in \tilde{\mathcal{W}}$ such that $X' \to \tilde{W} \to Y'$ and $H(\tilde{W}) < H(W')$. We have

$$\begin{aligned} p_{X,Y}(x,y) &= \sum_{w \in \mathcal{W}} p_W(w) p_{X|W}(x|w) p_{Y|W}(y|w) \\ &= \sum_{w \in \mathcal{W}'} p_W(w) p_{X|W}(x|w) p_{Y|W}(y|w) + \sum_{w \in \mathcal{W} \setminus \mathcal{W}'} p_W(w) p_{X|W}(x|w) p_{Y|W}(y|w) \\ &= \mathsf{P}\{W \in \mathcal{W}'\} p_{X',Y'}(x,y) + \sum_{w \in \mathcal{W} \setminus \mathcal{W}'} p_W(w) p_{X|W}(x|w) p_{Y|W}(y|w) \\ &= \mathsf{P}\{W \in \mathcal{W}'\} \sum_{w \in \tilde{\mathcal{W}}} p_{\tilde{W}}(w) p_{X'|\tilde{W}}(x|w) p_{Y'|\tilde{W}}(y|w) + \sum_{w \in \mathcal{W} \setminus \mathcal{W}'} p_W(w) p_{X|W}(x|w) p_{Y|W}(y|w). \end{aligned}$$

Thus, if we define

$$W'' = \begin{cases} W, & \text{if } W \in \mathcal{W}', \\ \tilde{W}, & \text{if } W \notin \mathcal{W}'. \end{cases}$$

Then $X \to W'' \to Y$ forms a Markov chain. Also

$$\begin{aligned} H(W'') &\overset{(a)}{=} H\left(\mathsf{P}\{W \in \mathcal{W}'\}\right) + \mathsf{P}\{W \in \mathcal{W}'\} H(W') + \mathsf{P}\{W \notin \mathcal{W}'\} H(W | W \notin \mathcal{W}') \\ &< H\left(\mathsf{P}\{W \in \mathcal{W}'\}\right) + \mathsf{P}\{W \in \mathcal{W}'\} H(\tilde{W}) + \mathsf{P}\{W \notin \mathcal{W}'\} H(W | W \notin \mathcal{W}') \\ &\overset{(b)}{=} H(W), \end{aligned}$$

where (a) and (b) follow since $H(X, \Theta) = H(\Theta) + H(X|\Theta)$. Thus we obtain $G(X;Y) < H(W)$, a contradiction.

Appendix G
Proof of Proposition 8

Proposition 6 implies $|\mathcal{W}| \le 3$.

Geometrically, the Markov chain $X \to W \to Y$ means that the vectors $\{p_{Y|X}(\cdot|x) : x \in \mathcal{X}\}$ lie in the convex hull of the vectors $\{p_{Y|W}(\cdot|w) : w \in \mathcal{W}\}$.

When we restrict $|\mathcal{W}| = 2$, an application of Lemma 2 together with the observation that $p_{W|X}(w|x) = 0 \iff p_{X|W}(x|w) = 0$ and the Markovity $X \to W \to Y$ immediately results in the two cases described in Proposition 8.

Similarly, when we let $|\mathcal{W}| = 3$, by Lemma 2, the supports of $p_{Y|W}(\cdot|w)$ are different for each $w \in \mathcal{W}$. Therefore, we have

$$p_{Y|W} = \begin{bmatrix} 1 & 0 & x \\ 0 & 1 & \bar{x} \end{bmatrix},$$

where each column corresponds to a value of W and $0 \leq x \leq 1$.

Similarly $p_{W|X}$ can be one of the 6 cases below. Each row corresponds to a W. Since there are 3 rows, there are $3! = 6$ cases.

$$p_{W|X} \in \left\{ \begin{bmatrix} y & z \\ \bar{y} & 0 \\ 0 & \bar{z} \end{bmatrix}, \begin{bmatrix} y & z \\ 0 & \bar{z} \\ \bar{y} & 0 \end{bmatrix}, \begin{bmatrix} 0 & \bar{z} \\ y & z \\ \bar{y} & 0 \end{bmatrix}, \begin{bmatrix} \bar{y} & 0 \\ y & z \\ 0 & \bar{z} \end{bmatrix}, \begin{bmatrix} \bar{y} & 0 \\ 0 & \bar{z} \\ y & z \end{bmatrix}, \begin{bmatrix} 0 & \bar{z} \\ \bar{y} & 0 \\ y & z \end{bmatrix} \right\},$$

for appropriately chosen numbers y, z between 0 and 1.

Assume

$$p_{Y|X} = \begin{bmatrix} \alpha & \beta \\ \bar{\alpha} & \bar{\beta} \end{bmatrix}$$

and $X \sim \text{Bern}(p)$. We first consider case 1 for $p_{W|X}$. In that case, we have the Markovity $p_{Y|W}p_{W|X} = p_{Y|X}$, i.e.,

$$\begin{bmatrix} 1 & 0 & x \\ 0 & 1 & \bar{x} \end{bmatrix} \begin{bmatrix} y & z \\ \bar{y} & 0 \\ 0 & \bar{z} \end{bmatrix} = \begin{bmatrix} \alpha & \beta \\ \bar{\alpha} & \bar{\beta} \end{bmatrix},$$

which yields $y = \alpha$ and $z + \bar{z}x = \beta$ or $z = \frac{\beta - x}{\bar{x}}$, hence $x \leq \beta$. Also

$$p_W = p_{W|X}p_X = \begin{bmatrix} py + \bar{p}z \\ p\bar{y} \\ \bar{p}\bar{z} \end{bmatrix}.$$

$y = \alpha$ is fixed, and z decreases monotonically with x. Since p_W is a linear function of z, the optimal z that minimizes the concave function $H(W)$ should lie at one of the end-points of z. The two end-points are $z = 0$ (corresponding to $x = \beta$) and $z = \beta$ (corresponding to $x = 0$). When $z = 0$, $H(W) > H(\bar{p}) = H(X)$ and may be eliminated because $H_1(X;Y) \leq H(X)$. When $z = \beta$, we have $x = 0$, hence $p_{Y|W}(y|1) = p_{y|W}(y|2)$ for all y.

In this case, we can replace W by it's sufficient statistic w.r.t Y, namely

$$T(W) = \begin{cases} 1, & \text{if w=0} \\ 0, & \text{otherwise.} \end{cases}$$

By property 6 from Section II this operation reduces the cardinality $|\mathcal{W}|$ and increases the entropy.

Cases $2, 3, 4$ for $p_{W|X}$ are very similar to case 1. Now consider case 5. In this case, Markovity implies

$$\begin{bmatrix} 1 & 0 & x \\ 0 & 1 & \bar{x} \end{bmatrix} \begin{bmatrix} \bar{y} & 0 \\ 0 & \bar{z} \\ y & z \end{bmatrix} = \begin{bmatrix} \alpha & \beta \\ \bar{\alpha} & \bar{\beta} \end{bmatrix}$$

and

$$p_W = \begin{bmatrix} p\bar{y} \\ \bar{p}\bar{z} \\ py + \bar{p}z. \end{bmatrix}$$

Thus we want to solve the following optimization problem:

$$\begin{aligned} &\text{Minimize } H(p\bar{y}, \bar{p}\bar{z}, py + \bar{p}z) \\ &\text{Subject to } \bar{y} + xy = \alpha, xz = \beta. \end{aligned}$$

We can eliminate variable x from the above constraints to obtain a constraint

$$\frac{\bar{\alpha}}{y} + \frac{\beta}{z} = 1.$$

We want to show that for the optimal (y^*, z^*), either $y^* = 1$ or $z^* = 1$. In that case, the support of W reduces to size 2, thus showing that $|\mathcal{W}| = 2$ suffices. To show this, we observe that $\{(y, z) : \frac{\bar{\alpha}}{y} + \frac{\beta}{z} = 1\}$ lies in the convex hull of the points $\{(1, \beta/\alpha), (\bar{\alpha}/\bar{\beta}, 1), (0, 0), (0, 1), (1, 0)\}$. Since $H()$ is a concave function, it's minimum in the convex hull should be one of the boundary points. It is easy to see that the minimum must be one of $\{(1, \beta/\alpha), (\bar{\alpha}/\bar{\beta}, 1)\}$. These 2 points satisfy the original constraint $\frac{\bar{\alpha}}{y} + \frac{\beta}{z} = 1$. Hence, either $y^* = 1$ or $z^* = 1$.

Case 6 is similar to case 5.

REFERENCES

[1] A. D. Wyner, "The common information of two dependent random variables," *IEEE Trans. Inf. Theory*, vol. 21, no. 2, pp. 163–179, Mar. 1975.

[2] R. Ahlswede and I. Csiszár, "Common randomness in information theory and cryptogra-phy—I: Secret sharing," *IEEE Trans. Inf. Theory*, vol. 39, no. 4, pp. 1121–1132, 1993.

[3] T. M. Cover, A. El Gamal, and M. Salehi, "Multiple access channels with arbitrarily correlated sources," *IEEE Trans. Inf. Theory*, vol. 26, no. 6, pp. 648–657, Nov. 1980.

[4] P. Cuff, "Distributed channel synthesis," *arXiv:1208.4415v3*, 2013.

[5] A. A. Borokov, "Mathematical statistics," *Gordon and Breach Science Publishers*, 1998.

[6] A. E. Gamal and Y. H. Kim, "Network information theory," *Cambridge University Press*, 2011.

[7] P. Gács and J. Körner, "Common information is far less than mutual information," *Probl. Control Inf. Theory*, vol. 2, no. 2, pp. 149–162, 1973.

[8] P. Cuff, "Communication in networks for coordinating behavior," *PhD Dissertation submitted to the dept. of Electrical Engg. at Stanford University*, 2009.

[9] P. Cuff, H. Permuter, and T. Cover, "Coordination capacity," *IEEE Transactions on Information Theory*, vol. 56, no. 9, p. 4181 4206, September 2010.

[10] K. Y, R. B, and C. V, "Matrix factorization techniques for recommender systems," *IEEE Computer Society*, 2009.

[11] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," *Advances in Neural Information Processing Systems 13: Proceedings of the 2000 Conference. MIT Press.*, pp. 556–562, 2001.

[12] J. Vant and R. Wilson, "A course in combinatorics," *Cambridge University Press*, 1992.